

Claims

We claim:

1. A method for preventing an unread activity from being bounced-back to an originating server during a replication operation, comprising:

storing an identification of an originating server of a replicated unread activity in an unread log of a receiving server; and

during a subsequent replication process initiated by the receiving server, preventing replication of the unread activity back to the originating server.

2. The method of claim 1, further comprising:

during the subsequent replication process, replicating the unread activity to at least one other server not identified as the originating server.

3. The method of claim 1, wherein storing an identification further comprises:

updating the unread log to include an unread entry corresponding to the replicated unread activity; and

storing the identification of the originating server with the unread entry.

4. The method of claim 3, wherein preventing the replication of the unread activity back to the originating server further comprises:

examining the unread log to determine if any unread entries stored therein correspond to an unread activity received from the originating server; and, during the subsequent replication process, not replicating any unread activity identified as being received from the originating server back to the originating server.

5. The method of claim 1, wherein the originating server has a name, and wherein the identification is a hash of the name of the originating server.

6. The method of claim 5, wherein during the subsequent replication process, if another server has the same hash as the originating server, the receiving server replicates the unread activity to the other server and back to the originating server.

7. The method of claim 6, wherein the originating server discards any duplicate replicated unread activities.

8. A bounce-back prevention system, comprising:

a receiving server for receiving an unread activity replicated by an originating server, the receiving server including an unread log for storing an identification of the originating server;
and

a system for preventing replication of the unread activity back to the originating server during a subsequent replication process initiated by the receiving server.

9. The system of claim 8, wherein the receiving server further comprises a replication system, and wherein the replication system of the receiving server replicates the unread activity to at least one other server not identified as the originating server during the subsequent replication process.

10. The system of claim 8, wherein the receiving server further comprises:

a system for updating the unread log to include an unread entry corresponding to the replicated unread activity, and for storing the identification of the originating server with the unread entry.

11. The system of claim 10, wherein the system for preventing the replication of the unread activity back to the originating server further comprises:

a system for examining the unread log to determine if any unread entries stored therein correspond to an unread activity received from the originating server; and
a system for preventing any unread activities, identified by the examining system as being received from the originating server, from being replicated back to the originating server, during the subsequent replication process.

12. The system of claim 8, wherein the originating server has a name, and wherein the identification is a hash of the name of the originating server.

13. The system of claim 12, wherein the receiving system includes a replication system, and wherein during the subsequent replication process, if another server has the same hash as the originating server, the replication system of the receiving server replicates the unread activity to the other server and back to the originating server.

14. The system of claim 13, wherein the originating server discards any duplicate replicated unread activities.

15. A program product stored on a recordable medium for preventing an unread activity from being bounced-back to an originating server during a replication operation, which when executed comprises:

program code for storing an identification of an originating server of a replicated unread activity in an unread log of a receiving server; and

program code for preventing replication of the unread activity back to the originating server, during a subsequent replication process initiated by the receiving server.

16. The program product of claim 15, further comprising:

program code for replicating the unread activity to at least one other server not identified as the originating server, during the subsequent replication process.

17. The program product of claim 15, wherein the program code for storing an identification further comprises:

program code for updating the unread log to include an unread entry corresponding to the replicated unread activity; and

program code for storing the identification of the originating server with the unread entry.

18. The program product of claim 17, wherein the program code for preventing the replication of the unread activity back to the originating server further comprises:

program code for examining the unread log to determine if any unread entries stored therein correspond to an unread activity received from the originating server; and,

program code for preventing replication of any unread activity the examining program code has

identified as being received from the originating server back to the originating server, during the subsequent replication process.

19. The program product of claim 15, wherein the originating server has a name, and wherein the identification is a hash of the name of the originating server.

20. The program product of claim 19, wherein the receiving server further includes program code for replicating the unread activity to the other server and back to the originating server during the subsequent replication process, if another server has the same hash as the originating server,

21. The program product of claim 20, wherein the originating server includes program code for discarding any duplicate replicated unread activities.

22. A method for preventing an unread activity from being bounced-back to at least one originating server during a replication operation, comprising:

- storing an identification of each originating server of a replicated unread activity in an unread log of a receiving server; and
- during a subsequent replication process initiated by the receiving server, preventing replication of the unread activity back to each originating server.